

Metaprogramming Applied to Numerical Problems

Mario Mulansky and Karsten Ahnert

Department of Physics and Astronomy, University of Potsdam

Metaprogramming is a programming technique where a computer program is written that writes or manipulates another program (or itself) that then solves a given task. This approach has several advantages over classical programming. Firstly, it leads to more efficient programming because the process of generating code is automated as you basically write code that writes code. Secondly, it might also lead to higher performance because using automatic code generation one can more easily produce code that can be optimized more efficiently by the compiler. A popular example is the Matrix Template Library (MTL) where Template Metaprogramming is used to create a high performance linear algebra library [1].

In C++ this is mostly realized by the use of templates and is then called *Template Metaprogramming* [2, 3]. Templates are placeholders for data types that are replaced – *instantiated* – by the compiler at compile time. Although templates were originally introduced to provide the functionality of parametrized types [4], it was quickly realized that they are indeed much more powerful and form a programming language themselves [5]. The first examples using C++ templates for Metaprogramming calculated prime numbers and were presented by E. Unruh on the C++ standardization meeting in 1994 in San Diego [6]. As most simple and instructive example serves the calculation of the factorial of a number at compile time [7]. Nowadays, Template Metaprogramming has found many applications, the most sophisticated are collected in the boost libraries, like Boost.MSM, Boost.Units, Boost.Accumulators and many others. In this work, we will use the ideas and techniques of Metaprogramming to implement numerical algorithms. This will heavily rely on existing Metaprogramming libraries, namely Boost.MPL [8] and Boost.Fusion [9] – compile time libraries that simplify the creation of Metaprogramming algorithms.

Metaprogramming is especially suitable to numerical algorithms where an approximate result is found by the successive application of a previously known number of stages where the type of computation at each state is similar. Perfect examples for such algorithms are the explicit Runge-Kutta schemes [10] for finding the solution of Ordinary Differential Equations (ODEs). In practice, solving ODEs numerically is usually done by finding the solution to an Initial Value Problem (IVP):

$$\dot{x}(t) = f(x, t), \quad x(t = 0) = x_0. \quad (1)$$

A Runge-Kutta scheme is defined by its number of stages s and its integration parameters $c_1 \dots c_s$, $a_{21}, a_{31}, a_{32}, \dots, a_{ss-1}$ and $b_1 \dots b_s$. The approximate solution for $x_1 \approx x(h)$ starting at x_0 is then found by

$$x_1 = x_0 + h \sum_{i=1}^s b_i F_i \quad \text{where} \quad F_i = f\left(x_0 + h \sum_{j=1}^{i-1} a_{ij} F_j, h c_i\right). \quad (2)$$

Depending on the number of stages s and the choice of parameters a, b, c , the approximate solution x_1 is exact up to some order p , which means $|x_1 - x(h)| \propto h^{p+1}$. The parameters a, b and c define the so-called Butcher tableau and fully describe the specific Runge-Kutta scheme. At each stage i the following calculations have to be performed:

$$F_i = f(y_i, h c_i), \quad y_{i+1} = x_0 + h \sum_{j=1}^i a_{i+1,j} F_j, \quad i = 1 \dots s-1 \quad \text{with} \quad y_1 = x_0 \quad (3)$$

$$F_s = f(y_s, h c_s), \quad x_1 = x_0 + h \sum_{j=1}^s b_j F_j. \quad (4)$$

In this work we introduce a generic C++ implementation of such a Runge-Kutta scheme based on Template Metaprogramming. Generic here means that our algorithms performs *any* Runge-Kutta scheme for which the user provides the Butcher array. Listing 1 shows the general usage of our generic implementation using the classical 4th order Runge-Kutta scheme [11] as an example. This scheme is the most common one and can be found in any textbook on numerical methods for ODEs, e.g. [12]. We start with defining the type representing the state of the system. As we here use the famous Lorenz system [13] as an introductory example, we use an array of doubles with length 3 as `state_type`. Before coming to the Runge-Kutta algorithm, the rhs of the ODE has to be defined. This is done by the function `lorenz` that calculates $f(x, t)$ for the Lorenz system writing the result into the variable `dxdt`. As a next step the parameters a, b, c for the Runge-Kutta scheme have to be defined, which is also done in terms of arrays of doubles. For simplicity, we here omit details on the types of the parameter arrays, like `coef_a_type`, but they are

Listing 1: Classical 4th order Runge-Kutta scheme realized using the generic algorithm. Note that some details are omitted here for simplicity, e.g. the coefficient types.

```
typedef array< double , 3 > state_type;

void lorenz( const state_type &x , state_type &dxdt , double t )
{ //defines the rhs of the ODE  $x' = f(x,t)$ , the result is written to dxdt
  dxdt[0] = sigma * ( x[1] - x[0] );
  dxdt[1] = R * x[0] - x[1] - x[0] * x[2];
  dxdt[2] = x[0]*x[1] - b * x[2];
}

int main() {

  // define the butcher array
  const array< double , 1 > a1 = {{ 0.5 }};
  const array< double , 2 > a2 = {{ 0.0 , 0.5 }};
  const array< double , 3 > a3 = {{ 0.0 , 0.0 , 1.0 }};

  const coef_a_type a = fusion::make_vector( a1 , a2 , a3 );
  const coef_b_type b = {{ 1.0/6.0 , 1.0/3.0 , 1.0/3.0 , 1.0/6.0 }};
  const coef_c_type c = {{ 0.0 , 0.5 , 0.5 , 1.0 }};

  // create the runge kutta scheme using the butcher array in a,b,c
  explicit_rk< state_type , 4 > rk4_stepper( a , b , c );

  // define the state  $x = (x,y,z)$  of the lorenz system
  state_type x = {{ 10.0 * rand() , 10.0 * rand() , 10.0 * rand() }};

  // make a number of timesteps
  for( size_t i=0 ; i<num_of_steps ; ++i, t+=dt )
    rk4_stepper.do_step( lorenz , x , t , dt );

  cout << x[0] << " , " << x[1] << " , " << x[2] << endl;
}
```

defined within our library and can be easily obtained. Finally, the Runge-Kutta stepper is created as an instance of the class `explicit_rk< state_type , 4 >`. Here, we tell the stepper as template parameters which state type to use (`array<double,3>`) and how many stages the scheme has (four). The constructor then takes the parameters defined above and we are ready to perform the integration by successive calls of the `rk4_stepper.do_step` function.

As seen in this example, with our library it is very easy to implement arbitrary Runge-Kutta schemes just by providing the corresponding Butcher arrays. The calculation details are hidden in the class `explicit_rk` which injects the given parameters into the algorithm and prevents the user from rewriting similar code when using different Runge-Kutta schemes. Of course, such a generic implementation can be obtained by a run-time implementation as well where basically a `for`-loop iterates through the stages. However, in such a case one finds a worse performance than for a direct implementation of the Runge-Kutta scheme as it is shown in our performance analysis. The reason is that in a run-time implementation the number of stages is unknown to the compiler which means that it has less possibilities for optimization.

The essential part of our algorithm is the class template `stage` representing one stage of the Runge-Kutta scheme, as shown in Listing 2. This class is meant to store the relevant information required to calculate one stage of the scheme, that is the number of this stage i , the parameter c_i and the parameters a_{ij} , b_j for the last stage respectively. Suppose we have an instance of such a stage filled with appropriate parameter values, then one stage of the according Runge-Kutta scheme can be calculated from Eqs. (3). Listing 3 shows the `calc_stage` function that does exactly this computation. The actual calculation is hidden in the function `vec_add`, which is implemented straight forward by a `for`-loop iterating through $j = 1 \dots i$. It is precisely at this point of the code where the strength of our approach comes into play: The stage number i of the stage to be currently calculated is known at compile time which allows for an efficient optimization of the code. Having a successive number of instances of the stage class with increasing stage number $i = 1 \dots s$ and filled with the parameters c_i , a_{ij} and b_i the complete Runge-Kutta step can be performed by successive calls of the `calc-stage`-function. The remaining task is to provide a convenient way to create and initialize

Listing 2: Definition of the fundamental stage class.

```
template< class T , size_t i >
struct stage // general (intermediate) stage
{
    T m_c; // parameter c_i
    boost::array< T , i > m_a; // parameters a_{i+1,1} ... a_{i+1,i}
                                // b_1 .. b_j respectively for the last stage
};
```

Listing 3: The `calc_stage` function performing one stage calculation.

```
state_type x; // the current state x
state_type x_tmp; // to store the intermediate result y_{i+1}
state_type F[stage_count]; // store the derivatives F_i
System system; // the rhs of the ODE, usually a function pointer

template< typename T , size_t i >
void calc_stage( const stage< T , i > &stage )
{ // performs one stage calculation of the i-th stage
    // obtain the rhs using x for the first stage, y_i otherwise
    if( i == 1 )
        system( x , F[i-1] , t + stage.c * dt );
    else
        system( x_tmp , F[i-1] , t + stage.c * dt );

    if( i < last_stage ) // performs the calculation x_tmp = x + dt * sum_j a_{ij}*F_j
        algebra<i>::vec_add( x_tmp , x , stage.a , F , dt);
    else // x = x + dt * sum_j b_j*F_j for the last, respectively
        algebra<i>::vec_add( x , x , stage.a , F , dt);
}
```

the stages and to perform the steps. Therefore, we implemented the `explicit_rk` class template which can be easily initialized and provides the `do_step` method for performing one complete Runge-Kutta step, as seen in Listing 1. Here, we made heavy use of the Boost.MPL and Boost.Fusion libraries which allowed us to write a sophisticated Meta-Algorithm in about 200 lines of code. Furthermore, the support of embedded Runge-Kutta methods which allow for error estimation is straight forward as it basically involves just the addition of another stage used to calculate the error estimate. This is also already implemented in our library and can be used by instantiating the `explicit_error_rk` class template.

The crucial point for us was to create a generic library that shows similar performance to direct implementations of the Runge-Kutta schemes. To check whether we have met this goal we did excessive performance tests comparing different implementations with the new generic approach. As reference implementation we used the routines in `odeint` [14] as they provide state-of-the art implementations of ODE solvers. We also adapted the algorithms from the “Numerical Recipes” (NR) and made use of the routines from the Gnu Scientific Library (GSL). Finally, we also implemented a runtime version of a generic Runge-Kutta without any Template Metaprogramming (rt gen, RK4 only). As actual integration schemes we first used the classical Runge-Kutta-4 algorithm, but also measured the runtime for the RK5(4) Cash-Karp scheme which involves error estimates. The measurements were performed for a number of modern machines and with different compilers, see table in Fig. 1. For each run we computed the relative performance (one over runtime) based on the runtime of `odeint`. The results are shown in Fig. 1. As seen there, `odeint`, our new generic implementation and the algorithms from the Numerical Recipes show basically the same performance, while the GSL and the runtime generic algorithm are significantly slower by about a factor 2. It should be noted that we found heavy fluctuation up to 50% between different compilers and in general the latest gcc version showed the best performance. However, we conclude from these tests that our generic Runge-Kutta algorithm based on Template Metaprogramming is indeed competitive in terms of run-time to direct implementations. This is a direct consequence of the compile-time programming as the run-time implementation is significantly slower.

As a summary, we here present a C++ implementation of a generic Runge-Kutta algorithm based on Template Metaprogramming. With generic we mean that our algorithm runs *any* explicit Runge-Kutta scheme where the user has only to provide the parameters a , b , c . The use of Template Metaprogramming allows for a run-time performance




FIG. 1: Performance of different Runge-Kutta implementations. The left panel shows the results for the classical RK4 scheme, while on the right the RK5(4) Cash-Karp scheme was used. The graphs show the performance relative to the `odeint`-implementation averaged over several runs on different machines using different compilers (see table on the right).

of the Runge-Kutta schemes comparable to, sometimes even faster than, a direct, straight forward implementation of the scheme. The advantage of our approach is that we completely separated the algorithm from the underlying parameters which reduces the effort of implementing new Runge-Kutta schemes to simple parameter initializations.

-
- [1] Jeremy Siek and Andrew Lumsdaine. The matrix template library. In Denis Caromel, Rodney Oldehoeft, and Marydell Tholburn, editors, *Computing in Object-Oriented Parallel Environments*, volume 1505 of *Lecture Notes in Computer Science*, pages 501–501. Springer Berlin, Heidelberg, 1998.
 - [2] A. Alexandrescu. *Modern C++ Design*. Addison-Wesley Professional, Boston, US-MA, 2001.
 - [3] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns*. Addison-Wesley, Boston, US-MA, 1994.
 - [4] B. Stroustrup. Parameterized types for C++. *Computing Systems*, 2(1), 1989.
 - [5] Todd L. Veldhuizen. C++ templates are turing complete. Technical report, 2003.
 - [6] E. Unruh. Prime number computation. ANSI X3J16-94-0075/ISO WG21-462, 1994.
 - [7] T. Veldhuizen. Using C++ template metaprograms. *C++ Report*, 7(4):36–43, 1995.
 - [8] D. Abrahams and A. Gurtovoy. *C++ Template Metaprogramming*. Addison-Wesley Professional, Boston, US-MA, 2004.
 - [9] J. de Guzman, D. Marsden, and T. Schwinger. *Fusion 2.0*.
 - [10] J. C. Butcher. *Numerical methods for ordinary differential equations*. John Wiley & Sons, Hoboken, US-NJ, 2003.
 - [11] W. Kutta. Beitrag zur näherungsweise Integration totaler Differentialgleichungen. *Z. für Math. Phys.*, 46:435–453, 1901.
 - [12] W. H. Press, S. T. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: the Art of Scientific Computing*. Cambridge University Press, Cambridge, England, second edition, 1992.
 - [13] E. N. Lorenz. Deterministic nonperiodic flow. *J. of the Atmospheric Sciences*, 20(2):130–141, 1963.
 - [14] K. Ahnert and M. Mulansky. *odeint* www.odeint.com.